IN THE CLAIMS:

This listing of claims replaces all prior versions and listings of claims in the application:

1-10.   (canceled)

11.   (original)   A method for managing identifiers allocated to one or more replicas of a database, the database comprising a data space comprising a plurality of data items, a global ID space comprising a plurality of identities (IDs) for identifying data items, the method comprising:

allocating a first range of IDs from the global ID space to a replica of the database;

receiving a request from the replica for a second range of IDs; and

providing a second range of IDs from the global ID space to the replica, a size of the second range of IDs being selected based upon a usage rate of IDs by the replica.

12.   (original)   The method of claim 11, wherein the size of the second range of IDs is selected based upon at least one of an average usage rate of IDs, a current usage rate of IDs, and a rate of change of usage rate of IDs over time by the replica.

13.   (original)   The method of claim 11, wherein the replica comprises a first replica, and wherein the first replica intermittently communicates with a second replica of the database for synchronizing data between the first and second replicas.

2

14.   (original)   The method of claim 13, wherein the size of the second range of

IDs is selected to provide sufficient numbers of IDs for the first replica to satisfy ID usage by the

first replica between successive communications with the second replica.

15.   (original)   The method of claim 14, wherein the second replica comprises a

master copy of the database.

16.   (original)   The method of claim 11, further comprising submitting a request

from the replica for the second range of IDs, and wherein the request is received by an ID

administrator associated the global ID space, the ID administrator allocating the second range of

IDs from the global ID space in a manner that prevents conflict with IDs allocated to another

replica.

17.   (original)   The method of claim 16, wherein the request for a second range of

IDs is submitted when a predetermined number of available IDs in the first range of IDs reaches

a trigger point.

18.   (original)   A method for managing identifiers allocated to a plurality of

replicas of a database comprising a data space including a plurality of data items, and an ID space

comprising a plurality of identities (IDs) for identifying data items, the method comprising:

3

allocating a first range of IDs n_1 from the ID space to a first replica of the database, the

IDs in the first range ranging from (x) to $(x+n_1-1)$, x being an integer;

allocating a second range of IDs n_2 from the ID space to a second replica of the database,

the IDs in the second range ranging from (y) to $(y+n_2-1)$, y being an integer greater than $(x+n_1-1)$;

receiving a request from a replica of the master database for a third range of IDs; and

providing the third range of IDs from the ID space to the requesting replica, a size n_3 of

the third range of IDs being selected based upon a usage rate of IDs by the requesting replica, the

IDs in the third range ranging from (z) to $(z+n_3-1)$, z being an integer greater than $(y+n_2-1)$.

19.   (original)   The method of claim 18, wherein the replica receiving the third

range of IDs is the first replica, and wherein the request is generated when a predetermined

number IDs from the first range of IDs remain unused by the first replica.

20-25. (canceled)

4